

REMARKS

Claims 1-20 remain pending in the application. Reconsideration is respectfully requested in light of the following remarks.

Section 112, First Paragraph, Rejection:

The Office Action rejected claims 1-20 under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement. More specifically, the Examiner asserts that the specification does not support “wherein the new data is new with respect to the production database”. Applicants respectfully submit that the specification clearly describes the relationship of the new data with respect to the production database. First, Applicants note that the Background section of the instant specification describes the problems associated with loading new data into production databases (e.g., data warehouses). The specification then goes on to describe how to allow new data to be loaded into the production database without interrupting access to the production database (e.g., by using the claimed methods and systems). One of skill in the art would clearly understand that the term “new data” refers to data that is new with respect to the production database. As one specific example, Applicants note paragraph [0036] which recites in part:

In one embodiment, a database clone may be created from a checkpoint of a production database...New data may be loaded into the database clone...During the refresh process, the production database is accessible by users. In one embodiment, once loading data and restructuring the database as necessary is complete, both the production database and the database clone may be shut down. The checkpoint may then be switched to the primary file system. The database clone may be switched to the production database. The production database may then be restarted. The production database has the latest, updated data. (Emphasis Added)

Thus, the specification clearly describes that a database clone is created and new data is loaded into the database clone as a refresh process of the production database. After the new data is loaded in the database clone and the checkpoint is switched to be the primary file system of the production database, the production database has the latest,

updated data. One of skill in the art can easily understand from this description (among numerous others in the specification) as well as the identified problem to be solved that the new data is new with respect to the production database. Furthermore, it is clear that the new data is not simply checkpoint data as asserted by the Examiner (addressed in more detail below). Accordingly, removal of the 112 rejection is requested.

Section 112, Second Paragraph, Rejection:

The Office Action rejected claims 1-20 under 35 U.S.C. § 112, second paragraph, as indefinite. More specifically, the Examiner asserts that the term “new” is a relative term which renders the claim indefinite. The Examiner further asserts that “new” is not defined by the claim, and the specification does not provide a standard for ascertaining the requisite degree, and one of ordinary skill in the art would not be reasonably apprised of the scope of the invention. Applicants strongly disagree with each of the Examiner’s assertions. First, the claim clearly defines the relationship of the new data with respect to the data of the production database. More specifically, the data is explicitly defined in the claims as new with respect to data of the production database. Applicants submit that it is clear (especially in light of the specification, similar to arguments above) that new production data is being loaded into the database clone and then the file system data of the production database is switched to the storage checkpoint which was just updated with the new data, thereby refreshing the production database. While new is a relative term, Applicants have provided a standard for ascertaining the requisite degree of newness by defining the relationship of the new data with respect to the production data in the claim. Applicants further submit that one of skill in the art could clearly understand the meaning of new data from the claimed limitations, especially in light of the present specification. Accordingly, removal of the 112 rejection is requested.

Section 103(a) Rejections:

The Examiner rejected claims 1, 2, 5-10, 13-16 and 19-20 under 35 U.S.C. § 103(a) as being unpatentable over Kampe et al. (U.S. Publication 2002/0032883)

(hereinafter “Kampe”) in view of Raman et al. (U.S. Publication 2003/0217119) (hereinafter “Raman”), claims 3, 11 and 17 as being unpatentable over Kampe in view of Raman and further in view of Ishihara et al. (U.S. Patent 6,636,876) (hereinafter “Ishihara”), and claims 4, 12 and 18 as being unpatentable over Kampe in view of Raman and further in view of AAPA. Applicants respectfully traverse these rejections for at least the following reasons.

Regarding claim 1, Kampe in view of Raman fails to teach or suggest loading new data to the database clone, wherein said load updates the storage checkpoint, wherein the new data is new with respect to the production database. With respect to this feature, the Examiner relies on Kampe. First, Applicants note that Kampe teaches a system for maintaining replicas of a primary component on nodes (via checkpointing) for achieving high availability in a cluster (see paragraph [0013]). Thus, Kampe is particularly directed towards achieving high availability of a component (in the case of failover) in a cluster rather than loading new data into a production database. For example, in Kampe, paragraph [0064] (which the Examiner relies on for switching from the previous file system data of the production database to the storage checkpoint to be the file system data for the production database), various failover scenarios are discussed where a primary component PC1 fails and the checkpoint is used to recreate the last consistent state of the primary component (either on the first node or on a second node). Thus, in Kampe, checkpoints are used to maintain the current state (or information for recovery) of a component on the first node, and are not used to load new data (with respect to the production database) into a production database using a database clone.

Applicants note that the Examiner specifically cites “(504)” for this limitation without any further explanation. 504 is a block in Figure 4A of Kampe which states “continuously update checkpoint”. Thus, 504 relates to maintaining the checkpoint with respect to the current state of the primary component. There is no indication in this Figure or anywhere else (in Kampe or Raman) which relates to loading new data to the database clone, wherein the new data is new with respect to the production database. Using checkpoints to maintain a current state (or backup of a current state) does not relate

to loading new information into a database clone (or checkpoint) which is new with respect to the production database (or, using the Examiner's citation, the primary component of Kampe). Thus, claim 1 relates to loading data, which did not previously exist in the production database, into a database clone / checkpoint, and then using the file system data of the checkpoint as the file system data of the production database (thereby loading the new data into the production database). Neither Kampe nor Raman (singly or in combination) teach these features of claim 1. Instead, as indicated above, Kampe maintains a current state of the primary component for failover purposes and does not load new information with respect to the primary component. Applicants further note that loading new data (as recited in Applicants' claims) into the checkpoints of Kampe would teach away from the purpose of high availability in a cluster since the replica / checkpoint would no longer represent the current state of the primary component. Thus, the cited references do not teach or suggest **loading new data to the database clone, wherein said load updates the storage checkpoint, wherein the new data is new with respect to the production database**.

In response to these arguments, the Examiner cites paragraph [0078] of Kampe and describes various data characteristics of the checkpoint of Kampe (e.g., format, states, control blocks, and attributes). The Examiner then goes on to assert these data characteristics "must be 'loaded' as claimed in order to manage the checkpoint". Applicants respectfully submit that this checkpoint attribute data does not relate to the new data of the claims. For example, the claimed limitations include generating a database clone which includes a storage checkpoint, loading new data to the database clone, which updates a storage checkpoint, which is used after the load as the file system data for the production database. Clearly, the new data is not simply characteristic data of the checkpoint (as asserted by the Examiner), but is instead new production data for the production database.

The Examiner goes on to assert that "Kampe also meets the claim limitation when interpreting 'new' as 'current', 'recent' or 'fresh'". The Examiner then goes on to refer to "continuously updating a checkpoint" and asserts that the continuously updated data

“thus recently comes into existence during creation/update in steps 503-505, whereas the data in the primary component has already existed”. Applicants submit that the interpretation of ‘new’ as already existing in the production database is entirely without merit. Clearly, if the data is the same data that has already existed in the primary component, it is not new data with respect to the production database (primary component in the Examiner’s citation). Applicants further submits that the previously submitted arguments from above appear to be ignored in these assertions.

Furthermore, the Examiner has failed to provide a proper motivation or reason to combine to the two references. In the instant Office Action, the Examiner asserts it would have been obvious to modify Kampe, “to improve accessibility for file system data, as known to one of ordinary skill in the art and taught by Raman (para. 0049)”. The Examiner has not provided any motivation or reason to make the specific combination and instead has merely provided a portion of Raman which describes a method for providing copies of consistent file systems with concurrent read-write updating of the file system. There is no indication in this section (or in the Examiner’s provided motivation) as to why the maintaining of file systems as taught by Raman applies to the “primary component” of Kampe. More specifically, the cited portion does not relate to a “production database”, the accessibility of a production database during the load of new data, or any reason to include such features in Kampe as alleged by the Examiner. Instead, Raman teaches that updates may be made over an IP network, and that concurrent read-only access to the remote copies is available.

In response to these arguments, the Examiner asserts “one would have used Raman to provide access to file system data and availability during loads (see below) with the motivation of increasing the flexibility of Kampe by allowing Kampe to operate with files system data and to have availability during loads”. Thus, the Examiner only identifies the presumed result of the combination and does not identify any particular reason for one of skill in the art to combine the references. Additionally, combining references only to achieve the result of the claims is no more than hindsight reasoning and is improper.

Thus, for at least the reasons provided above, Applicant submits that Kampe and Raman, taken singly or in combination, fail to teach all the features and limitations of claim 1, and so Applicant submits that claim 1 and those claims dependent therefrom are patentably distinct and non-obvious over the cited art, and are thus allowable.

Claims 8, 9, and 15 include similar limitations as claim 1, and so the above arguments apply with equal force to these claims. Thus, for at least the reasons provided above, Applicant submits that claims 8, 9, and 15, and those claims respectively dependent therefrom, are patentably distinct and non-obvious, and are thus allowable.

Applicants also assert that numerous ones of the dependent claims recite further distinctions over the cited art. However, since the rejections have been shown to be unsupported for the independent claims, a further discussion of the dependent claims is not necessary at this time.

CONCLUSION

Applicants submit the application is in condition for allowance, and notice to that effect is respectfully requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-12400/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Applicants

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8850

Date: April 28, 2008